arXiv:1612.01397v1 [cs.LG] 5 Dec 2016

Implicit Modeling – A Generalization of Discriminative and Generative Approaches

Dmitrij Schlesinger, Carsten Rother*

Dresden University of Technology

Abstract

We propose a new modeling approach that is a generalization of generative and discriminative models. The core idea is to use an implicit parameterization of a joint probability distribution by specifying only the conditional distributions. The proposed scheme combines the advantages of both worlds – it can use powerful complex discriminative models as its parts, having at the same time better generalization capabilities. We thoroughly evaluate the proposed method for a simple classification task with artificial data and illustrate its advantages for real-word scenarios on a semantic image segmentation problem.

1 Introduction

We start with an overview that illustrates the proposed modeling approach in a broader context. In statistical modeling it is assumed that *there exists* a generating probability distribution that generates data and the goal is to search for a modeling probability distribution that fits the data adequately. Even if it is not necessary to model the generating probability distribution completely (e.g. in case of discriminative models, see later) we prefer to start the considerations with a joint probability distribution $p(x, y)$, where x is an observation and y is an unknown associated label that we want to obtain, e.g. the class label in a classification task.

In order to predict y from x, a common approach is to first model the joint distribution $p(x, y)$, in particular by decomposing it as $p(x, y){=}p(x|y){\cdot}p(y)$. This is called generative modeling, where a prior over the labels $p(y)$, and a likelihood $p(x|y)$ are specified. Prediction is then carried out by employing $p(y|x){\propto}p(x|y){\cdot}p(y)$, obtained via Bayes' rule. Generative modeling often admits natural decompositions of the joint distribution: for example, in image segmentation $p(y)$ is a prior model over "good" segmentations, and $p(x|y)$ describes the image formation process. The learning of the unknown parameters

*This work was supported by German Federal Ministry of Education and Research (BMBF, 01IS14014A-D) and European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 647769). The computations were performed on an HPC Cluster at the Center for Information Services and High Performance Computing (ZIH) at TU Dresden.

is often carried out by maximizing the joint log-likelihood of the training data

$$\arg\max_{\theta_1,\theta_2} \sum_{t=1}^{T} \Big[\log p(y_t;\theta_1) + \log p(x_t|y_t;\theta_2) \Big], \tag{1}$$

where $\big((x_t, y_t), t{=}1 \ldots T\big)$ is the training sample and θ_1 and θ_2 are unknown parameters to be learned.

Alternatively, one can decompose $p(x, y){=}p(y|x){\cdot}p(x)$. While this kind of decomposition is often less intuitive, it has the property that $p(x)$ can be arbitrary and thus be ignored, since only $p(y|x)$ is required to predict y from x. This is called discriminative modeling, where only the conditional $p(y|x)$ is specified. A prominent example of such a modeling approach are conditional random fields (CRF) [6]. The statistical learning is usually done by maximizing the conditional log-likelihood

$$\arg\max_{\theta} \sum_{t=1}^{T} \log p(y_t|x_t;\theta). \tag{2}$$

Discriminative models often exhibit excellent performance when training data is abundant. Besides they are less vulnerable to misspecification (see e.g. [11]) – i.e. the situation that the "true" generating probability distribution is not contained in the modeled family. However, discriminative models are prone to over-fitting when the parameters of the distribution $p(y|x)$ are learned from *little data*.

Related works. In order to alleviate the vulnerability of discriminative models to over-fitting, regularization in form of prior knowledge about the parameters is often used. In fact, the unknown parameters of the probability distribution are handled as additional random variables. Consequently, regularization to prevent over-fitting is often done in an ad-hoc manner by imposing simple mathematically convenient priors on model parameters in form of e.g. sparsity constraints, quadratic penalties, etc. (see e.g. [9] for an overview of related methods). Usually, the assumptions about the used parameter prior are far from being well founded (e.g. Gaussian in the parameter space). Moreover, the design of such a prior is usually less intuitive as compared with the design of the original model.

Another line of work (e.g. [2, 8]) proposes to combine generative and discriminative likelihood functions for learning. An example is [8], where the so-called multi-conditional log-likelihood $\log\big(p(y|x;\theta)^{\alpha} p(x|y;\theta)^{\beta}\big)$ is maximized during training. Here two conditional models share parameters, and α, β are used to weight the contributions of the two conditionals. Note that the combination is done on the algorithmic level, rather than on the modeling one. Basically, the proposed multi-conditional log-likelihood does not correspond to any probability and hence is not statistically well-motivated.

In [1] a learning objective is defined by using a pair of models each with its own set of parameters. Blending between generative and discriminative models is achieved by putting a prior on the two sets of parameters: fully independent parameters lead to a fully discriminative model, and constraining parameters to be identical recovers the fully generative model. Note that the method operates with *two different* models. Furthermore, the used probability distributions have to belong to the *same* family, since their parameters must be comparable.

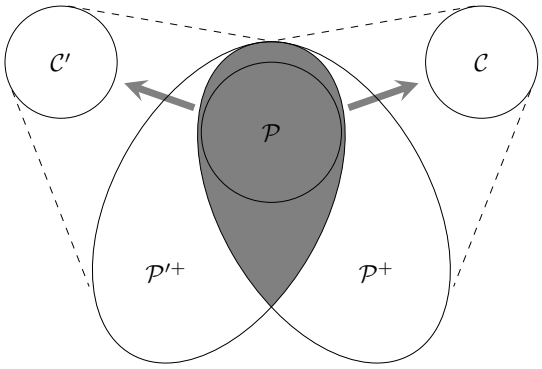

Figure 1: The considered families of probability distributions (see text for explanations). The implicitly defined family of probability distributions of interest is shaded by gray color.

Contribution. We propose a modeling framework that has the following properties. First of all, it is a generalization of commonly used generative and discriminative modeling approaches. Blending between generative and discriminative extremes is achieved by appropriate design of the corresponding model parts. The framework is statistically well motivated – the learning procedure searches for a joint probability distribution in a well (although not explicitly) defined family. We have no prior for the parameter (in contrast to e.g. [9]). There is no model approximation as in [8]. In contrast to e.g. [1] we are not restricted in the choice of the conditional distributions. We can combine state-of-the-art discriminative models $p(y|x)$ with arbitrary generative components for $p(x|y)$, increasing thereby the generalization capabilities of the former.

2 Implicit Models

From now on we denote random variables as well as their ranges by capitalized symbols, i.e. X for observations and Y for hidden variables. Particular values are denoted by $x \in X$ and $y \in Y$. Similarly, $\theta \in \Theta$ denotes a value of an unknown parameter, where Θ is the parameter space. Consequently, $p(X, Y; \theta)$ denotes a particular joint probability distribution, whereas $p(X, Y; \Theta)$ is a family of probability distributions.

Let us recall the main advantages of generative and discriminative models, namely: (i) the generative models are more robust, i.e. less vulnerable to overfitting, (ii) the discriminative models are more powerful and less vulnerable to misspecification for large data. The reason for such a behavior can be explained as follows. Consider a family of the joint probability distributions $\mathcal{P}{=}p(X, Y; \Theta)$ (see Fig. 1). Furthermore, consider the corresponding family of conditional (posterior) probability distributions $\mathcal{C}{=}p(Y|X; \Theta_1)$, i.e. those which can be derived from members of $\mathcal{P}$. For example, let $\mathcal{P}$ be a family of Gaussian mixture models with two Gaussians having the same variance, i.e. $p(x, y){=}p(y){\cdot}p(y|x)$, with the hidden variable $Y = \{1, 2\}$ and $p(x|y; \mu, \sigma) = \mathcal{N}(x; \mu_y, \sigma)$. Then $\mathcal{C}$ consists of all possible linear logistic regressions.

Note however, that at the same time there exists a corresponding family of prior distributions for observations $p(X;\Theta)$, i.e. all those which can be obtained by marginalization of members of $\mathcal{P}$ over Y. The crucial assumption of the discriminative modeling is that the prior $p(X)$ can be arbitrary, which implicitly enlarges the family $\mathcal{P}$. Let us consider another family $\mathcal{P}^+$ of joint probability distributions, namely all those, whose conditional distributions are members of $\mathcal{C}$. For the above example with Gaussians, $\mathcal{P}^+$ will consists of all joint probability distributions that have a linear logistic regression as the posterior. Obviously, the family $\mathcal{P}^+$ is a superset of $\mathcal{P}$, i.e. in $\mathcal{P}$ the family of priors $p(X;\Theta)$ is specified (i.e. restricted), whereas in $\mathcal{P}^+$ the priors are arbitrary. Generative models work with $\mathcal{P}$ that is usually too restricted. It leads to worse recognition performance and vulnerability to model misspecification but better generalization capabilities. Discriminative models work in fact with $\mathcal{P}^+$ that is too general. Note that $\mathcal{P}^+$ can be understood as the cartesian product $\mathcal{C}\times p(X;\Theta)$, where $\Theta = \mathbb{R}^{|X|}$. Hence, $\mathcal{P}^+$ is extremely huge. If labeled training data is abundant, learning discriminative models is nevertheless able to "recover" the true generating probability distribution, which may be the member of $\mathcal{P}^+$ but not of $\mathcal{P}$. Hence, employing discriminative models alleviates the misspecification problem. However, it suffers from over-fitting for little data.

Our goal now is to develop an approach that allows to specify families of the joint probability distributions that are larger than $\mathcal{P}$ but more restricted than $\mathcal{P}^+$. The core idea is based on the following observation. In addition to the families $\mathcal{P}$, $\mathcal{C}$ and $\mathcal{P}^+$, consider the corresponding family of conditional probability distributions for observations given labels $\mathcal{C}' = p(X|Y;\Theta_2)$ – i.e. those which can be derived from members of $\mathcal{P}$. Furthermore, consider the family of joint probability distributions $\mathcal{P}'^+$ whose conditionals are members of $\mathcal{C}'$. Again, the family $\mathcal{P}'^+$ is a superset of $\mathcal{P}$. Hence, the intersection $\mathcal{P}^+ \cap \mathcal{P}'^+$ is also a superset of $\mathcal{P}$. At the same time, this intersection is a subset of both $\mathcal{P}^+$ and $\mathcal{P}'^+$. Our goal is to learn joint models that are in this intersection. To summarize, we propose to *implicitly* model the joint distribution $p(X,Y)$ by means of the two conditional distributions $p(Y|X)$ and $p(X|Y)$.

Definition 1 (Implicit Modeling)**.** *Let two families $p(Y|X;\Theta_1)$ and $p(X|Y;\Theta_2)$ of conditional probability distributions be given. A joint probability distribution $p(X,Y)$ belongs to the family $p(X,Y;\Theta_1,\Theta_2)$ of interest if there exist $p(X)$, $p(Y)$, $\theta_1 \in \Theta_1$ and $\theta_2 \in \Theta_2$ satisfying*

$$p(x) \cdot p(y|x;\theta_1) = p(y) \cdot p(x|y;\theta_2) \quad \forall x \in X, y \in Y. \tag{3}$$

Given that, the joint probability distribution is defined by $p(x,y;\theta_1,\theta_2) = p(x) \cdot p(y|x;\theta_1)$.

Obviously, the above conditions imply that the families $p(X)$ and $p(Y)$ of marginal probability distributions are not arbitrary but implicitly restricted by Θ_1, Θ_2 through (3). Hence, in further we will write $p(X;\Theta_1,\Theta_2)$ having in mind that $p(X)$ should satisfy (3) (for $p(Y;\Theta_1,\Theta_2)$ analogously).

Weak implicit modeling

In further we weaken conditions (3). The main motivation to do that is the following. In practice, we would not like "to care" about the modeling of the joint

probability distribution at all. In other words, we would not like to *derive* the families of conditionals from a common family of joint probability distributions. We would like just to *design* two families of conditional probability distributions, hence implicitly defining the family of the corresponding joint ones. Moreover, we would like to be able to combine *nearly arbitrary* families of conditionals, that need not to have something in common from the modeling point of view. For example, we may want to use elaborated discriminative model (like e.g. complex Conditional Random Fields or Convolutional Neural Networks) for $p(Y|X;\Theta_1)$, combining them with physically motivated forward models (e.g. some rendering engines that produce images x given a scene description y) for $p(X|Y;\Theta_2)$. However, given two particular conditionals $p(Y|X;\theta_1)$ and $p(X|Y;\theta_2)$, it may be the case that there is no priors $p(X)$ and $p(Y)$ satisfying (3). Moreover, it may be even the case that given two families $p(Y|X;\Theta_1)$ and $p(X|Y;\Theta_2)$, there is no pair of θ_1 and θ_2 for which there exist $p(X)$ and $p(Y)$ satisfying (3). Consider the following example. Let one of the families of conditionals, let say $p(X|Y;\Theta_2)$, is designed in such a way, that it leads to independent joint probability distribution, i.e. $p(x|y;\theta_2) = p(x|y';\theta_2)$ for all x, all pairs y and y' and all θ_2. Let the conditionals from the other family do not have this property. Obviously, according to the first family, the implicitly defined joint probability distribution should be independent, whereas according to the second one, it is not. Consequently, for such a pair of families there is no joint probability distribution satisfying (3).

Note that the conditions (3) are posed for the marginal probabilities, but not for the joint ones. Hence, we weaken (3) as follows:

Definition 2 (Weak Implicit Modeling)**.** *Let two families $p(Y|X;\Theta_1)$ and $p(X|Y;\Theta_2)$ of conditional probability distributions be given. A joint probability distribution $p(X,Y)$ belongs to the family $p(X,Y;\Theta_1,\Theta_2)$ of interest if there exist $p(X)$, $p(Y)$, $\theta_1 \in \Theta_1$ and $\theta_2 \in \Theta_2$ satisfying*

$$\begin{aligned} p(y) &= \sum_{x\in X} p(x)\cdot p(y|x;\theta_1), \quad \forall\ y \in Y, \\ p(x) &= \sum_{y\in Y} p(y)\cdot p(x|y;\theta_2), \quad \forall\ x \in X. \end{aligned} \tag{4}$$

Given that, the joint probability distribution is defined either by $p(x,y;\theta_1,\theta_2)$=$p(x)\cdot p(y|x;\theta_1)$ or by $p(x,y;\theta_1,\theta_2)$=$p(y)\cdot p(x|y;\theta_2)$ depending on the application.

Obviously, (4) are necessary but not sufficient conditions for (3). The following lemma asserts the applicability of the weak implicit modeling for situations described above, namely if the considered families of conditional probability distributions are not derived from a common joint model, but designed in an application specific way independently from each other.

Lemma 1. *Under some mild conditions there exists a pair of marginal probability distributions $p(X)$ and $p(Y)$ satisfying (4) for any pair of conditional probability distributions $p(Y|X;\theta_1)$ and $p(X|Y;\theta_2)$ (see proof in Appendix A).*

It is also easy to see that conditions (4) implicitly define *two* (possibly) different joint probability distributions: one being defined by $p(x,y;\theta_1,\theta_2) =$

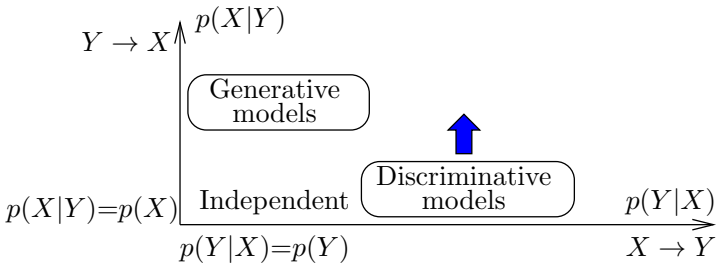

Figure 2: A common framework of generative, discriminative and weak implicit modeling approaches. Axes correspond to the weakness of the corresponding parts. "$p(Y|X) = p(Y)$" means that $p(Y|X)$ does not depend on X, "$X \to Y$" means a deterministic mapping. For $p(X|Y)$ analogously.

$p(x;\theta_1,\theta_2) \cdot p(y|x;\theta_1)$ and the other defined by $p(x,y;\theta_1,\theta_2) = p(y;\theta_1,\theta_2) \cdot p(x|y;\theta_2)$. In practice, one should decide which of these two is of interest. In most cases we are interested in recognition at the end, for which only the posterior $p(Y|X)$ is necessary for inference. Therefore it seems reasonable to define the family of joint probability distributions so that its induced posterior coincides with the conditional $p(Y|X;\theta_1)$ that we model. Hence, in further we define the joint by $p(x,y;\theta_1,\theta_2) = p(x;\theta_1,\theta_2) \cdot p(y|x;\theta_1)$. Of course, in other application scenarios it might be the other case.

We would like to point in that although we weakened the original conditions (3), the family of joint probability distributions $p(x,y;\theta_1,\theta_2) = p(x;\theta_1,\theta_2) \cdot p(y|x;\theta_1)$ satisfying (4) is still considerably restricted as compared to the standard discriminative approach. Remember that in the latter the marginal $p(X)$ is completely free, whereas in the former the set of $p(X)$ is restricted by (4).

The proposed modeling approach is neither generative nor discriminative (in the meaning described at the very beginning). It is closer to a generative one, in spirit, since we are working with joint probability distributions (although not explicitly defined). At the same time it is easy to see that both fully generative and discriminative models are special cases of the weak implicit modeling, by choosing the families of the conditional distributions as follows. When $p(X|Y;\Theta_2)$ does not depend on Y, i.e. $p(x|y)$=$p(x|y')$ for any x, y, y' and θ_2, we obtain a fully discriminative model. Similarly, we are in the generative extreme if $p(Y|X;\Theta_1)$ does not depend on X. On the other hand, if $p(Y|X;\Theta_1)$ is deterministic, i.e. it is a mapping $X \to Y$, we deal with discriminant functions. Hence, the whole spectrum of modeling approaches fits into a common framework illustrated in Fig. 2. The axes correspond to the "weakness" of the constituents. Commonly used fully discriminative and fully generative approaches occupy relatively small areas along the axes. We attempt to explore the area in the middle, among other things, to push the discriminative approaches towards the direction depicted by the blue arrow in the figure, increasing their robustness to overfitting.

The "weakness" of the used conditional probability distributions (and hence the position of an approach in the proposed framework in Fig. 2) can be controlled by choosing how strongly for example $p(X|Y)$ depends on Y. In practice it can be done e.g. in the following manner. Consider a distribution defined as

$p(x|y) \propto \exp(\alpha E(x,y) + E(x))$ with some energy functions $E(\cdot)$. Setting α to zero leads to independence, a strong α corresponds to a deterministic mapping, i.e. there is a (unique) value of y for each x. If necessary, such a controlling factor can be set manually, weighting the "weakness/importance" of the corresponding family. On the other hand, it can also be considered as a usual unknown parameter and learned from data.

3 Learning Implicit Models

Given a training sample of observation-label pairs $\big((x_t, y_t), t{=}1 \ldots T\big)$, we aim to learn our implicit model by maximizing the joint log-likelihood

$$\begin{aligned} &\arg\max_{\theta_1,\theta_2} \sum_{t=1}^{T} \log p(x_t, y_t; \theta_1, \theta_2) = \\ &\qquad = \arg\max_{\theta_1,\theta_2} \sum_{t=1}^{T} \Big[\log p(x_t; \theta_1, \theta_2) + \log p(y_t|x_t; \theta_1)\Big]. \end{aligned} \tag{5}$$

For gradient-based optimization, we need to differentiate (5) wrt. θ_1, θ_2. It is straightforward to do that for the conditional $p(y|x;\theta_1)$. In contrast, the differentiation of the first addend in (5) is less trivial because there is no closed-form expression for it, i.e. the marginal distribution $p(X;\theta_1,\theta_2)$ depends *implicitly* on both θ_1 and θ_2.

We model $p(X;\theta_1,\theta_2)$ by means of a generation process that samples observations according to the desired probability distribution. Let us assume an (infinite) Markov chain that generates sequences $(x^0, y^0, x^1, y^1, \ldots, x^{n-1}, y^{n-1}, x^n)$, $n{\to}\infty$ by sampling from the corresponding conditional distributions $p(Y|X;\theta_1)$ and $p(Y|X;\theta_2)$ respectively, i.e. y^i is drawn from $p(Y|X{=}x^i;\theta_1)$ and x^i is drawn from $p(X|Y{=}y^{i-1};\theta_2)$. Obviously, stationary distribution of such a Markov chain satisfies (4) (it follows directly from Lemma 1, see Corollary in Appendix A), i.e. our $p(X;\theta_1,\theta_2)$ is the stationary distribution of this Markov chain. Hence, we model $p(x_t;\theta_1,\theta_2)$ with the probability that a Markov chain as described above generates the training example x_t.

For better readability we omit here the detailed derivation of the algorithm for optimizing (5). More detailed explanations can be found in appendix B. In short: the probability $p(x_t;\theta_1,\theta_2)$ is obtained by marginalization over all sequences $(x^0, y^0, x^1, y^1, \ldots, x^{n-1}, y^{n-1}, x^n{=}x_t)$ generated by the Markov chain considered above. This summation is infeasible especially if x and y are complex by themselves, e.g. images and labelings (the state space of the Markov chain is of huge dimensionality). Hence, we use stochastic gradient ascent for approximation – instead to marginalize over all sequences, we draw just one from the Markov chain. To summarize, we need to carry out two steps to perform stochastic gradient ascent for optimizing (5) (on the example of just one training example (x^*, y^*)): (i) Generate a sequence of x-s and y-s starting from the training example x^* according to the probability distributions $p(Y|X;\theta_1)$ and $p(X|Y;\theta_2)$ given the current model parameters, and (ii) consider certain generated pairs as the "additional labeled training examples" for computing gradients wrt. θ_1 and θ_2. One stochastic gradient ascent step for one training example (x^*, y^*) is summarized in Algorithm 1 and illustrated in Fig. 3 (for many exam-

Algorithm 1 Stochastic gradient ascent step for (5) (one training example)

Input: Training example $(x^*, y^*) \equiv (x_t, y_t)$,
current parameters θ_1, θ_2
Output: Updated model parameters θ_1', θ_2'
Sample $\tilde{y}$ according to $p(Y|X{=}x^*; \theta_1)$
Sample $\tilde{x}$ according to $p(X|Y{=}\tilde{y}; \theta_2)$
Sample $\hat{y}$ according to $p(Y|X{=}\tilde{x}; \theta_1)$
Compute sufficient statistics η for the generated pairs
Compute gradients and do steepest ascent with a small step size λ:

$$\theta_1' = \theta_1 + \lambda \cdot \big(\eta_1(\tilde{x}, \tilde{y}) - \eta_1(\tilde{x}, \hat{y}) + \eta_1(x^*, y^*) - \eta_1(x^*, \tilde{y})\big)$$
$$\theta_2' = \theta_2 + \lambda \cdot \big(\eta_2(x^*, \tilde{y}) - \eta_2(\tilde{x}, \tilde{y})\big)$$

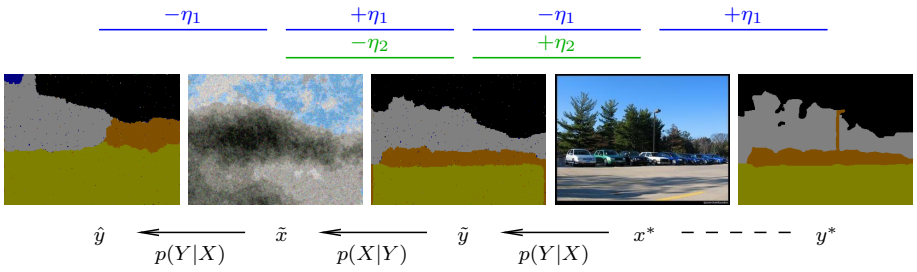

Figure 3: Illustration of a single step of stochastic gradient ascent. The examples in the chain are generated in the direction shown by the black arrows using the current conditional probability distributions. The pairs marked blue are used to learn $p(Y|X; \theta_1)$. The gradient is thereby obtained as the difference of sufficient statistics for the corresponding pairs. The pairs marked green are used to learn $p(X|Y; \theta_2)$. Images are taken from our experiment with semantic segmentation (see the next section). More qualitative results can be found in Appendix C.

ples their gradients should be averaged). Thereby, we assume that the conditional probability distributions of interest are both members of the exponential family, i.e. they can be written in form e.g. $p(y|x; \theta_1) \propto \exp\big(\langle \eta_1(x, y), \theta_1 \rangle\big)$, where $\eta_1(x, y)$ are sufficient statistics. Hence, the necessary gradients are just differences of the corresponding sufficient statistics.

Interestingly, the algorithm is somewhat similar to the standard conditional likelihood learning. More precisely, if we omit the sampled $\tilde{x}$ and $\hat{y}$ and use only the pairs (x^*, y^*) and $(x^*, \tilde{y})$ for updating the gradient, it would be exactly the conditional likelihood. The crucial difference here is that not only pairs from the training sample are used for learning $p(Y|X; \theta_1)$ but also pairs that are generated by the current conditional probability distributions $p(Y|X; \theta_1)$ and $p(X|Y; \theta_2)$, which implicitly define the current joint probability distribution $p(X, Y, \theta_1, \theta_2)$.

At this point we would like to note that our algorithm reminds to some extent on the Contrastive Divergence [5], since we also generate chains of a finite length starting from training examples in order to get examples for gradient cal-

culation. We would like to point in that our algorithm is by no means a variant of Contrastive Divergence, because the latter just performs another task. It is designed to draw examples from a *joint* probability distribution given *explicitly*. It is done by alternate sampling from the conditional probability distributions that are *derived* from the target joint one (i.e. Gibbs Sampling). Usually, both conditional probability distributions are simply given by means of an energy function that is *the same* for both conditionals. Given a generated chain, the *last* example is assumed to be drawn according to the needed joint probability distribution and is used for gradient calculation. In contrast, our algorithm does not need single examples drawn from the joint probability distribution, but the whole chains that are generated according to the certain rules (see Appendix B for details). Thereby *all* generated examples are used for gradient calculation. Moreover, we do not specify the joint probability distribution at all. Only conditional probability distributions are given, that are not derived from a common joint one but just belong to pre-defined families, that might have nothing in common (no similar structure, no common parameters, etc.).

4 Experiments

Classification example with artificial data. We start with an artificial example to illustrate our modeling approach. First, we define a generating joint probability distribution, from which training and test samples are drawn. Thus we define it in a generative manner, i.e. $p(x, y) = p(y) \cdot p(x|y)$. The variable Y is discrete and can take three values $y \in \{1, 2, 3\}$ (we will call them "classes"), the prior probability distribution of classes is uniform. The observation X is a real number, the conditional probability distribution $p(X|Y)$ consists of Gaussians (one per class) of the same variance but with different mean values. To summarize, the generating model is

$$p(x, y) = \frac{1}{3} \cdot \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{(x - \mu_y)^2}{2\sigma^2}\right]. \tag{6}$$

In particular, we use the variance $\sigma = 1$ and mean values $\mu_y = -1, 0, 1$ for three classes in this experiment.

The learned implicit model consists of two parts. The first one (we refer it as discriminative part) is the conditional probability distribution $p(Y|X)$ of classes given an observation, for which we choose quadratic logistic regression $p(y|x) \propto \exp[a_y \cdot x^2 + b_y \cdot x + c_y]$, $y \in \{1, 2, 3\}$. We intentionally use a model that may over-fit (linear logistic regression would be sufficient for the true generating model as described above). The second part (called the generative one) is the conditional probability distribution for observations given classes, for which we use one Gaussian per class. To be consistent with the above notations we use the exponential family like parameterization $p(x|y) \propto \exp[d_y \cdot x^2 + e_y \cdot x]$. Note that this model may over-fit as well, as we allow the Gaussians to have different d_y that correspond to different variances.

We perform several experiments with an increasing size of training samples to analyze the generalization capabilities of the models and learning approaches that we consider. Each experiment consists of the following: first, a training sample of a particular size is drawn from the generating probability distribution. After that learning is performed, where we compare the following schemes:

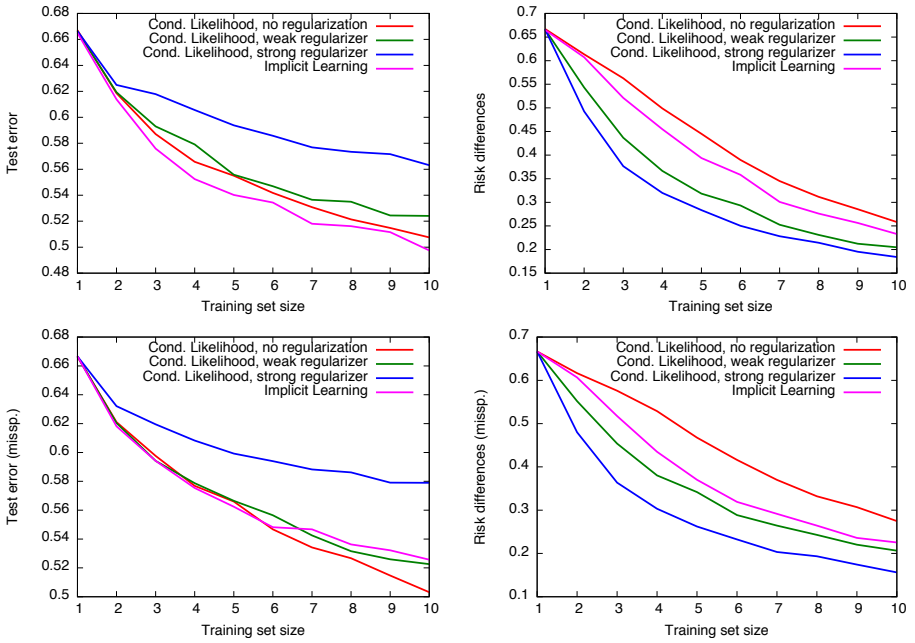

Figure 4: Results for experiments with artificial data. Top row – experiment without misspecification, bottom row – with misspecification. Left: test error rates, right: risk differences.

1. **Standard conditional likelihood**, which we consider as the baseline. In addition, we perform learning with a quadratic regularizer for model parameters – using "weak" (a small weighting constant) and "strong" regularization.

2. **Learning of the implicit model** as described in the previous section (we refer it as "implicit learning").

For the learned model we define the classifier to be the Maximum A-posteriori decision and compute two measures:

1. **Test error rate.** The (average) error rate on a very large test sample, i.e. the Bayesian risk for the learned classifier.

2. **Risk difference.** Additionally, we compute the absolute value of the difference between the error rate on the training sample (the empirical risk) and the error rate on the test sample (the Bayesian risk, when the test sample is large enough). We would like to emphasize that this measure is crucial to adequately assess the generalization capabilities of the different models and learning strategies, because it represents a kind of "guarantees" for predicting unseen data. Smaller risk difference means thereby better generalization capability.

For each particular size of the training set we repeat the experiment several times and report average measurements. Fig. 4 (top row) shows the dependencies of the considered measures on the size of the training set. First of all, one

can clearly see that the implicit learning consistently outperforms the standard conditional likelihood approach (without regularization) with respect to both measures. With respect to the recognition rate the implicit learning is also consistently better as compared to the regularized likelihood (for both weak and strong regularization). The most robust learning is the conditional likelihood with strong regularization. It is easy to explain by the fact that both the empirical and the Bayesian risks in this situation are worse as compared to other methods (see the blue line in Fig. 4 (top, left)). Hence, the usual trade-off "performance vs. stability" is clearly seen.

The focus in the previous experiment was mainly to analyze the behavior of the various schemes with respect to over-fitting and generalization. There is also another (to some extent "opposite") problem that is often crucial in many real applications, namely the problem of misspecification. The situation is that the true (generating) probability distribution is not contained in the family of modeling probability distributions. It is well known that discriminative modeling performs better for larger training samples (in comparison to generative modeling), because the set of modeling probability distributions is larger. The aim of the following experiment is to study the influence of the introduced generative part on the behavior of the learning with respect to misspecification for smaller training samples. To address this problem we simulate misspecification in the following way. First, we restrict the generative part $p(X|Y)$ of our implicit model to be Gaussians with the *same* variance. At the same time, in the generating model we use Gaussians of *different* variances. Note that the posterior $p(y|x) \propto \exp[a_y \cdot x^2 + b_y \cdot x + c_y]$ alone is still well-specified. The generative part $p(X|Y)$ is not indeed. Consequently, the family of joint probability distributions that are modeled implicitly by the pair $p(X|Y)$ and $p(Y|X)$ is misspecified as well. In other words, the usage of such a misspecified $p(X|Y)$ considerably restricts the family of modeling joint probability distributions. We perform the same kind of experiments as done previously, the results are shown in Fig. 4 (bottom row). As expected, the implicit learning performs slightly worse with respect to the recognition rate as compared to the case without misspecification. At the same time, the discriminative learning is better for larger training samples. Nevertheless, implicit learning still slightly outperforms the baseline for smaller samples. Moreover, the stability (risk difference) is almost not affected, hence the generalization capabilities of the implicit learning is still considerably superior.

Semantic Segmentation. In the next experiment we study the behavior of implicit modeling for the application of semantic image segmentation. The task is to label each pixel of an image by a class from a predefined label-set L. Formally, the hidden vector y is a labeling $y : V \to L$ of a graph $G = (V, E)$ whose nodes $i \in V$ correspond to image pixels and whose edges $(i, j) \in E$ encode a neighborhood structure (we use 8-neighborhood). The observation x is an image $x : V \to C$, i.e. mapping that assigns a color $c \in C$ to every pixel of the image. Among other things, in this experiment we would like to simulate a situation that is very common in Computer Vision. Assume that there exists already a model and we want to improve it using it just as a "black box", i.e. to combine it with some other model or to build something on top of it. We use pixel-wise independent decision forest classifier as such a "black

box"[1]. Let $z_i(x) \in L$ be the output label predicted by the decision forest for the pixel i. Based on that we define the following conditional random field to specify $p(Y|X)$:

$$p(y|x) \propto \exp\Big[\sum_{i\in V} q\big(y_i, z_i(x)\big) + \sum_{ij\in E} \big(a_t(y_i, y_j) + b_t(y_i, y_j) \cdot ||x_i - x_j||^2\big)\Big] \quad (7)$$

The first part of the energy represents unary potentials. The table $q : L\times L \to \mathbb{R}$ reflects the "reliability" of the decision forest's output. The pairwise potentials are linear functions of the squared color differences $a + b \cdot ||x_i - x_j||^2$ in the corresponding neighbouring pixels. The coefficients a and b depend on edge type (horizontal, vertical or diagonal, referred by index t in (7)) as well as on the label pair (y_i, y_j) on this edge. To summarize, given a decision forest we have $9 \times |L|^2$ additional free parameters in our CRF.

For the generative part (i.e. the probability distribution $p(X|Y)$) we use a model that is similar in spirit to a Gaussian mixture model. To this end we introduce latent variables g_i for each pixel i that represent "color numbers". In addition the coloring is required to be smooth inside the segments. To summarize, the model is

$$p(x, g|y) \propto \exp\Big[\sum_{i\in V}\big(h(y_i, g_i) + c \cdot ||x_i||^2 + \\ + \langle d(g_i), x_i\rangle\big) + e\cdot \sum_{ij\in E} \delta(y_i{=}y_j)\cdot||x_i - x_j||^2\Big]. \quad (8)$$

The weight matrix h assigns a value for each pair (label, color number). The coefficient c (shared for all colors) and color specific vectors d control the distribution of RGB-values given the color number. The constant e controls the smoothness of coloring inside the segments. $\delta(\cdot)$ is the Kronecker delta – it gives 1 if its argument is true. The necessary $p(x|y)$ is obtained by marginalizing (8) over g. All unknown parameters are learned in our experiments.

Recall that for learning we need to draw samples from $p(Y|X)$ and $p(X|Y)$, that is not a trivial task by itself, since both (7) and (8) are CRF-s. We use Gibbs Sampling for this purpose. In addition, in order to accelerate the overall learning procedure, we use "warm start" at each gradient update iteration as follows. Let e.g. $\tilde{y}_t$ denotes the example $\tilde{y}$ sampled from the Markov chain for the training example $(x_t, y_t) = (x^*, y^*)$ (see Alg. 1). At the beginning of the learning procedure we initialize the necessary labeling examples from all chains (i.e. $\tilde{y}_t$, $\hat{y}_t$ for all training examples (x_t, y_t)) randomly. The coloring examples $\tilde{x}_t$ are initialized by the corresponding original images x_t. At each gradient update iteration of the algorithm we draw new examples by Gibbs Sampling, starting from examples obtained in the previous gradient update iteration. For example, the new $\tilde{y}_t$ is obtained by Gibbs Sampling using the current conditional probability distribution $p(Y|X = x_t; \theta_1)$, starting from the old $\tilde{y}_t$ (from the previous iteration), performing thereby only a small number of Gibbs Sampling iterations. In practice, just one Gibbs Sampling iteration per gradient update usually performs best, taking into account both the quality of the results and computational demand. Similar acceleration tricks are widely used in the literature (see e.g. [12, 3]).

[1] We used the code from [10].

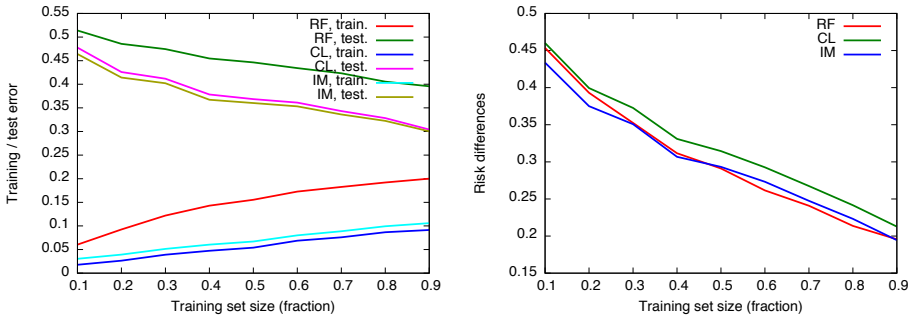

Figure 5: Results for semantic segmentation. RF – pixel-wise independent Random Forest classifier. CL – Conditional Likelihood learning. IM – implicit model. Left: training and test errors. Right: error differences.

We use the Stanford Background Dataset [4] for this experiment. Due to the computational demand of Gibbs Sampling, we do not use the whole dataset and only choose 100 images randomly. The training examples (up to 90 images) are chosen randomly out of these 100, the rest is used for testing. For each training sample we first learn a pixel-independent Decision Forest classifier using standard methods. For the corresponding CRF (7) we again compare two different scenarios: standard conditional likelihood learning and implicit learning. After the models are learnt we use the maximum marginal decision strategy for inference. The "recognition error" is the relative number of misclassified pixels, i.e. the Hamming distance averaged over the training/test sample. The measures of interest are as before the recognition error itself and the averaged absolute difference between the training and the test errors.

The results are presented in Fig. 5 (some qualitative results are shown in appendix C). Despite the test error we give also the training error for clarity. In addition, the rates for basic Random Forest classifier are also given. The typical behaviour can be observed. Building a CRF on top of the Random Forest improves both training and test errors. However, the training error is improved more essentially as compared to the test error. Hence, the stability becomes worse (compare the red and the green lines in Fig. 5, right). It is indeed expected, because in the CRF additional free parameters are introduced. Using implicit modeling makes the *training* error *worse* as compared to the Conditional Likelihood. However, at the same time the *test* error becomes *better*. Hence, by using implicit modeling we get the stability back. To summarize, the implicit learning outperforms the standard Conditional Likelihood with respect to both recognition rate and stability for the whole range of training sample sizes.

5 Conclusion

We presented a new modeling approach called implicit modeling. The families of joint probability distributions of interest are modeled implicitly by specifying two families of conditional probability distributions. We showed that the weak implicit modeling is a generalization of both generative and discriminative approaches. Being able to use powerful discriminative models as its constituents,

the method has at the same time better generalization capabilities. Experiments on both artificial and real data confirm the statement.

There are numerous directions for further works. The presented approach can work with nearly arbitrary constituents. Hence, an interesting research direction would be to apply the method to complex discriminative models such as decision tree fields [10] or convolutional neural networks [7].

In this paper we mainly focused on modeling aspects and used maximum likelihood for learning. Other choices are possible as well. For instance, using pseudo-likelihood may lead to more efficient learning. It would be also interesting to apply other learning objectives like loss-driven learning or composite likelihood to implicitly defined models.

In our opinion, an essential advantage of the implicit modeling in comparison to standard discriminative models is the ability (at least theoretically) to learn unknown parameters in a fully unsupervised manner (since the *joint* likelihood is optimized). It gives the possibility to train complex posterior probability distributions using incomplete or very weakly labeled data, that is often highly desirable in practice. We did not elaborate this direction so far, hence, it will be one of our further works.

References

[1] Christopher M Bishop and Julia Lasserre. Generative or discriminative? Getting the best of both worlds. *Bayesian Statistics*, 8:3–23, 2007.

[2] Guillaume Bouchard and Bill Triggs. The trade-off between generative and discriminative classifiers. In *IASC International Symposium on Computational Statistics*, pages 721–728, 2004.

[3] L.-C. Chen, A. G. Schwing, A. L. Yuille, and R. Urtasun. Learning deep structured models. In *ICML*, pages 1785–1794, 2015.

[4] S. Gould, R. Fulton, and D. Koller. Decomposing a scene into geometric and semantically consistent regions. In *Computer Vision, 2009 IEEE 12th International Conference on*, pages 1–8, Sept 2009.

[5] Geoffrey E. Hinton. Training products of experts by minimizing contrastive divergence. *Neural Comput.*, 14(8):1771–1800, August 2002.

[6] John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of the Eighteenth International Conference on Machine Learning*, ICML '01, pages 282–289, 2001.

[7] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. *CVPR (to appear)*, November 2015.

[8] Andrew McCallum, Chris Pal, Greg Druck, and Xuerui Wang. Multi-conditional learning: Generative/discriminative training for clustering and classification. In *AAAI*, 2006.

[9] Sebastian Nowozin and Christoph H. Lampert. Structured learning and prediction in computer vision. *Found. Trends. Comput. Graph. Vis.*, 6(3–8211;4):185–365, March 2011.

[10] Sebastian Nowozin, Carsten Rother, Shai Bagon, Toby Sharp, Bangpeng Yao, and Pushmeet Kohli. Decision tree fields. In *ICCV*, 2011.

[11] Patrick Pletscher, Sebastian Nowozin, Pushmeet Kohli, and Carsten Rother. Putting map back on the map. In *DAGM*, 2011.

[12] T. Tieleman. Training restricted boltzmann machines using approximations to the likelihood gradient. In *ICML*. ACM, 2008.

A Proof of Lemma 1

Let us first recall the definitions and the statement.

Definition 1 (Implicit Modeling)**.** *Let two families $p(Y|X;\Theta_1)$ and $p(X|Y;\Theta_2)$ of conditional probability distributions be given. A joint probability distribution $p(X,Y)$ belongs to the family $p(X,Y;\Theta_1,\Theta_2)$ of interest if there exist $p(X)$, $p(Y)$, $\theta_1 \in \Theta_1$ and $\theta_2 \in \Theta_2$ satisfying*

$$p(x) \cdot p(y|x;\theta_1) = p(y) \cdot p(x|y;\theta_2) \quad \forall x \in X, y \in Y. \tag{9}$$

Given that, the joint probability distribution is defined by $p(x,y;\theta_1,\theta_2) = p(x) \cdot p(y|x;\theta_1)$.

Definition 2 (Weak Implicit Modeling)**.** *Let two families $p(Y|X;\Theta_1)$ and $p(X|Y;\Theta_2)$ of conditional probability distributions be given. A joint probability distribution $p(X,Y)$ belongs to the family $p(X,Y;\Theta_1,\Theta_2)$ of interest if there exist $p(X)$, $p(Y)$, $\theta_1 \in \Theta_1$ and $\theta_2 \in \Theta_2$ satisfying*

$$\begin{aligned} p(y) &= \sum_{x \in X} p(x) \cdot p(y|x;\theta_1), \quad \forall\ y \in Y, \\ p(x) &= \sum_{y \in Y} p(y) \cdot p(x|y;\theta_2), \quad \forall\ x \in X. \end{aligned} \tag{10}$$

Given that, the joint probability distribution is defined either by $p(x,y;\theta_1,\theta_2){=}p(x){\cdot}p(y|x;\theta_1)$ or by $p(x,y;\theta_1,\theta_2){=}p(y){\cdot}p(x|y;\theta_2)$ depending on the application.

Lemma 1. *Under some mild conditions there exists a pair of marginal probability distributions $p(X)$ and $p(Y)$ satisfying (10) for any pair of conditional probability distributions $p(Y|X;\theta_1)$ and $p(X|Y;\theta_2)$.*

Proof. Let us represent the probability distribution $p(X)$ by a vector $v_x{\in}\mathbb{R}^{|X|}$ and the probability distribution $p(Y)$ by a vector $v_y{\in}\mathbb{R}^{|Y|}$. Similarly, $p(Y|X)$ can be understood as a matrix A of size $|Y|{\times}|X|$, where matrix elements a_{ij} satisfy $a_{ij} \geq 0$ and $\sum_i a_{ij} = 1$ for all j. Analogously, $p(X|Y)$ is another matrix B of $|X|{\times}|Y|$ elements. Hence, conditions (10) can be rewritten as

$$v_y = A \cdot v_x, \quad v_x = B \cdot v_y. \tag{11}$$

Let us substitute the first part of (11) in the second one and obtain

$$v_x = B \cdot A \cdot v_x = C \cdot v_x. \tag{12}$$

Note that both A and B are stochastic matrices. Hence, so is their product $C{=}B{\cdot}A$ as well. It means that under some mild conditions (for example, it is enough that all elements of C are strictly positive that is almost always the case in practice) there exists v_x satisfying (12), which is the eigenvector of C to the eigenvalue 1. To summarize, let us define v_x to be the eigenvector of C, and v_y as $v_y = A \cdot v_x$. Obviously, all conditions (10)-(12) are satisfied.

□

Corollary. Let us consider the Matrix $C{=}B{\cdot}A$ as a matrix of transition probabilities of a Markov chain, that generates sequences of x-s $(x^0, x^1, \ldots, x^n)$, $n{\to}\infty$. Obviously, the above considered vector v_x (and hence, the probability distribution $p(X)$) is the stationary distribution of this Markov chain. On the other hand, let us consider now the matrix $D{=}A{\cdot}B$ that can be understood as the matrix of transition probabilities for a Markov chain, that generates sequences of y-s $(y^0, y^1, \ldots, y^n)$, $n{\to}\infty$. Obviously, the eigenvector of D also satisfies all conditions (10)-(12), and hence, is also the stationary distribution of the corresponding Markov chain. Let us now consider a Markov chain that generates sequences $(x^0, y^0, x^1, y^1, \ldots, x^{n-1}, y^{n-1}, x^n, y^n)$, $n{\to}\infty$ by alternate sampling from the corresponding distributions $p(Y|X)$ and $p(Y|X)$ respectively, i.e. y^i is drawn from $p(Y|X{=}x^i)$ and x^i is drawn from $p(X|Y{=}y^{i-1})$. It is easy to see that both x^n and y^n obey the corresponding stationary distributions $p(X)$ and $p(Y)$ respectively.

B Algorithm for Learning Implicit Models

Here we give a more detailed derivation of the algorithm for learning implicit models. First, let us recall the problem to be solved. Given a training sample of observation-label pairs $\big((x_t, y_t), t{=}1\ldots T\big)$, we aim to learn our implicit model by maximizing the joint log-likelihood

$$\begin{aligned}\arg\max_{\theta_1,\theta_2} \sum_{t=1}^{T} \log p(x_t, y_t; \theta_1, \theta_2) = \\ = \arg\max_{\theta_1,\theta_2} \sum_{t=1}^{T} \Big[\log p(x_t; \theta_1, \theta_2) + \log p(y_t|x_t; \theta_1)\Big]. \end{aligned} \tag{13}$$

For gradient-based optimization, we need to differentiate (13) wrt. θ_1, θ_2. It is straightforward to do that for the conditional $p(y|x; \theta_1)$. In contrast, the differentiation of the first addend in (13) is less trivial because there is no closed-form expression for it, i.e. the marginal distribution $p(X; \theta_1, \theta_2)$ depends *implicitly* on both θ_1 and θ_2.

We model $p(X; \theta_1, \theta_2)$ by means of a generation process that samples observations according to the desired probability distribution. Let us assume an (infinite) Markov chain that generates sequences $(x^0, y^0, x^1, y^1, \ldots, x^{n-1}, y^{n-1}, x^n)$, $n{\to}\infty$ by sampling from the corresponding conditional distributions $p(Y|X; \theta_1)$ and $p(Y|X; \theta_2)$ respectively, i.e. y^i is drawn from $p(Y|X{=}x^i; \theta_1)$ and x^i is drawn from $p(X|Y{=}y^{i-1}; \theta_2)$. Obviously, stationary distribution of such a Markov chain satisfies (10) (see the Corollary above), i.e. our $p(X; \theta_1, \theta_2)$ is the stationary distribution of this Markov chain. Hence, we model $p(x_t; \theta_1, \theta_2)$ with the probability that a Markov chain as described above generates the training example x_t.

To ease exposition, let us assume for now that we want to maximize the probability of a single observation x^* under $p(X; \theta_1, \theta_2)$, i.e. only the first part of (13) for just one example. Consider a sequence

$$(z, x^*) = (x^0, y^0, x^1, y^1, \ldots, x^{n-1}, y^{n-1}, x^*) \tag{14}$$

generated by the Markov chain whose last entry is the observation x^* (z summarizes all generated x-s and y-s but x^*, Z denotes the set of all sequences).

The probability of this sequence is

$$p(z, x^*; \theta_1, \theta_2) = p(x^0) \cdot p(y^0|x^0; \theta_1) \cdot \\ \cdot \left[\prod_{i=1}^{n-1} p(x^i|y^{i-1}; \theta_2) \cdot p(y^i|x^i; \theta_1) \right] \cdot p(x^*|y^{n-1}; \theta_2). \tag{15}$$

The probability of the marginal for the observation x^* is obtained by marginalization over all possible sequences that end with x^*:

$$p(x^*; \theta_1, \theta_2) = \sum_z p(z, x^*; \theta_1, \theta_2). \tag{16}$$

Note that this probability does not depend on x_0 in (15) if the chain has infinite length, and can be neglected in practice if it is run long enough. Note that now $p(x^*)$ depends on the unknown parameters θ_1 and θ_2 in an *explicit* manner through the generating Markov chain. Hence, we can compute the gradient of the log marginal

$$\begin{aligned} \frac{\partial}{\partial \theta_{1,2}} \log p(x^*; \theta_1, \theta_2) &= \frac{\partial}{\partial \theta_{1,2}} \log \sum_z p(z, x^*; \theta_1, \theta_2) = \\ &= \left[\sum_{z'} p(z', x^*; \theta_1, \theta_2) \right]^{-1} \cdot \sum_z \frac{\partial p(z, x^*; \theta_1, \theta_2)}{\partial \theta_{1,2}} = \\ &= \sum_z \frac{p(z, x^*; \theta_1, \theta_2)}{p(x^*; \theta_1, \theta_2)} \cdot \frac{\partial \log p(z, x^*; \theta_1, \theta_2)}{\partial \theta_{1,2}} = \\ &= \sum_z p(z|x^*; \theta_1, \theta_2) \cdot \frac{\partial \log p(z, x^*; \theta_1, \theta_2)}{\partial \theta_{1,2}} \end{aligned} \tag{17}$$

that we require for learning (13). In practice, marginalization over all sequences z in (17) is intractable, especially when x and y are complex (e.g. images or labelings). Note however that the gradient in (17) is an expectation over the probability distribution $p(Z|X^n{=}x^*; \theta_1, \theta_2)$. Hence, we can use *stochastic gradient ascent*, i.e. we replace the expectation of a random variable by its realization. To summarize, we need to carry out two steps to perform stochastic gradient ascent for (17):

1. Generate a sequence $\hat{z}$ according to the probability distribution $p(Z|X^n{=}x^*; \theta_1, \theta_2)$ given the current model parameters, and
2. compute the gradient of $\log p(\hat{z}, x^*; \theta_1, \theta_2)$ with respect to θ_1, θ_2.

Concerning the second step, we obtain

$$\begin{aligned} \frac{\partial \log p(\hat{z}, x^*; \theta_1, \theta_2)}{\partial \theta_1} &= \sum_{i=0}^{n-1} \frac{\partial \log p(\hat{y}^i|\hat{x}^i; \theta_1)}{\partial \theta_1}, \\ \frac{\partial \log p(\hat{z}, x^*; \theta_1, \theta_2)}{\partial \theta_2} &= \sum_{i=1}^{n-1} \frac{\partial \log p(\hat{x}^i|\hat{y}^{i-1}; \theta_2)}{\partial \theta_2} + \\ &\quad + \frac{\partial \log p(x^*|\hat{y}^{n-1}; \theta_2)}{\partial \theta_2}, \end{aligned} \tag{18}$$

where $\hat{x}^i$ and $\hat{y}^i$ are elements of the generated $\hat{z}$. This can be interpreted as follows. After a sequence $\hat{z}$ is generated, its pairs $(\hat{x}^i, \hat{y}^i)$ at "odd" positions can be considered as an "additional labeled training sample", for which the conditional log-likelihood of $p(Y|X;\theta_1)$ has to be maximized with respect to θ_1. The same holds for "even" positions of the chain (including the transition from $\hat{y}^{n-1}$ to x^*) and the conditional probability distribution $p(X|Y;\theta_2)$ wrt. θ_2.

Unfortunately, generating a sequence $\hat{z}$ for the first step is not as easy as just to generate a sequence from a Markov chain. This is because we need to generate sequences according to $p(Z|X^n{=}x^*;\theta_1,\theta_2)$, i.e. conditioned by the last chain member. Common approaches for that would be e.g. importance sampling, sampling with rejection or Gibbs sampling. Unfortunately, these methods usually turn out to be too time-consuming, especially for complex x and y. Therefore in practice, we assume that our Markov chain is "reversible" in the sense that we can generate the sequence in the "opposite" direction, i.e. starting from x^* towards x^0, with the same probability. Note that this assumption is true, if the current conditional probability distributions $p(Y|X;\theta_1)$ and $p(X|Y;\theta_2)$ lead to existence of a *unique* joint probability distribution $p(X,Y;\theta_1,\theta_2)$, i.e. if the conditions (9) hold and we are in the case of (non-weak) implicit modeling. Otherwise, we assume that the two joint probability distributions induced by (10) (see the definition of the weak implicit modeling) are close enough to each other. Hence, examples drawn from $p(x,y;\theta_1,\theta_2){=}p(x){\cdot}p(y|x;\theta_1)$ and from $p(x,y;\theta_1,\theta_2){=}p(y){\cdot}p(x|y;\theta_2)$ obey nearly the same joint probability distribution. Observe that in doing so the pairs $(\hat{x}^i, \hat{y}^i)$ from the sequence $\hat{z}$ that were generated using $p(X|Y;\theta_2)$ (during the generation in opposite direction) serve as the training data to learn $p(Y|X;\theta_1)$ and vice versa.

Exponential family. In order to do learning as described above, the conditional probabilities $p(y|x;\theta_1)$ and $p(x|y;\theta_2)$ must be differentiable wrt. θ_1 and θ_2 respectively. Furthermore, we should be able to draw samples from both distributions. To make this slightly more concrete, let us assume that the conditional distributions both belong to the exponential family, i.e. they can be written in the form

$$\begin{aligned} p(y|x;\theta_1) &= \frac{1}{Z_1(x,\theta_1)} \cdot \exp\big(\langle \eta_1(x,y), \theta_1 \rangle\big), \\ p(x|y;\theta_2) &= \frac{1}{Z_2(y,\theta_2)} \cdot \exp\big(\langle \eta_2(x,y), \theta_2 \rangle\big), \end{aligned} \tag{19}$$

where $Z_{1,2}$ are the partition functions, $\eta_{1,2}$ are the sufficient statistics, and $\theta_{1,2}$ the unknown parameters; $\langle \cdot,\cdot \rangle$ denotes the inner product between two vectors. For a particular pair (x,y) the gradients of the conditional log-likelihoods are

$$\begin{aligned} \frac{\partial \log p(y|x;\theta_1)}{\partial \theta_1} &= \eta_1(x,y) - \frac{\partial \log Z_1(x,\theta_1)}{\partial \theta_1} = \\ &= \eta_1(x,y) - \mathbb{E}_{p(Y|X;\theta_1)}\big[\eta_1(x,y)\big], \\ \\ \frac{\partial \log p(x|y;\theta_2)}{\partial \theta_2} &= \eta_2(x,y) - \frac{\partial \log Z_2(y,\theta_2)}{\partial \theta_2} = \\ &= \eta_2(x,y) - \mathbb{E}_{p(X|Y;\theta_2)}\big[\eta_2(x,y)\big], \end{aligned} \tag{20}$$

where $\mathbb{E}_p[\cdot]$ denotes the expectation of a random variable $[\cdot]$ over a probability distribution p. Similar to the marginalization over all sequences z (see above), we use stochastic approximation, i.e. we exchange the expectation of a random variable by its realization. Hence, the stochastic gradient for $p(y|x;\theta_1)$ in (20) is obtained as the difference of sufficient statistics $\eta_1(x,y) - \eta_1(x,\hat{y})$ with an example $\hat{y}$ drawn from the probability distribution $p(Y|X;\theta_1)$ (for $p(X|Y;\theta_2)$ analogously). Note that we already need to draw these examples in order to generate the chain $(\hat{z}, x^*)$, so we can use them for gradient calculation (20), i.e. we do not need to extra generate them.

So far we focused on maximization of $p(x^*)$, i.e. only the first addend in (13). Obviously, we should take into account the second addend in (13) as well. Its stochastic gradient however is again the difference of the corresponding sufficient statistics (20). So we have only to add this gradient part to the previous one.

To conclude, one stochastic gradient ascent step for one training example (x_t, y_t) is summarized in Algorithm 1 and illustrated in Fig. 3 (see the main part of the paper). It is easy to see that the influence of the generated pairs on the log-likelihood decreases with the distance between the example position and the position of x^* in the chain. In practice we find that it is enough to generate a chain of minimal length, as required to compute the necessary sufficient statistics.

C Qualitative Results for Semantic Segmentation

Here we would like just to give an impression about qualitative results of our approach applied for semantic segmentation. In Figs. 6 and 7 chains of examples generated during the learning of implicit models are presented (see also Fig. 3 in the main part of the paper). The last column is the max-marginal decision for the corresponding image.

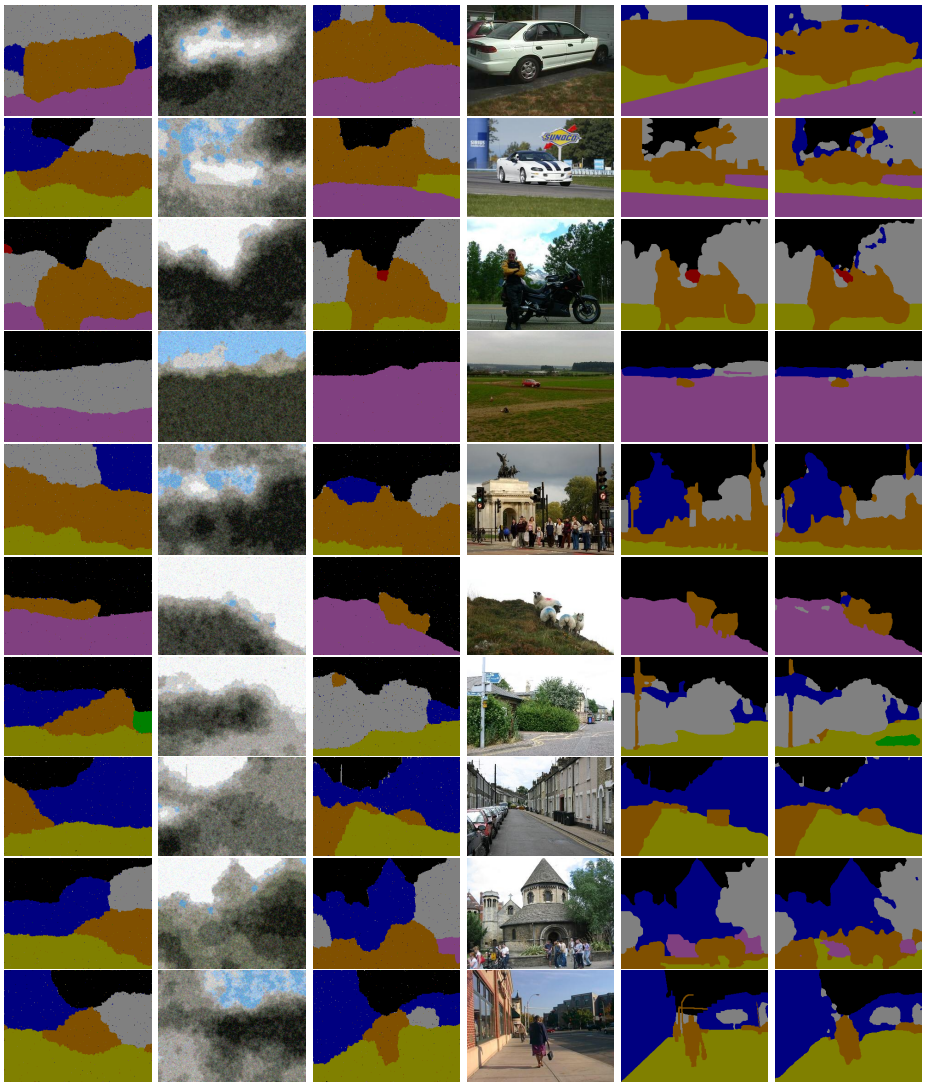

Figure 6: Chains of examples generated during the learning of implicit models. The first three columns are $\hat{y}$, $\tilde{x}$ and $\tilde{y}$ respectively (see Alg. 1). The fourth and fifth columns are samples from the training set, i.e. the observation x^* and the ground truth y^*. The last column is the max-marginal decision for the corresponding image.

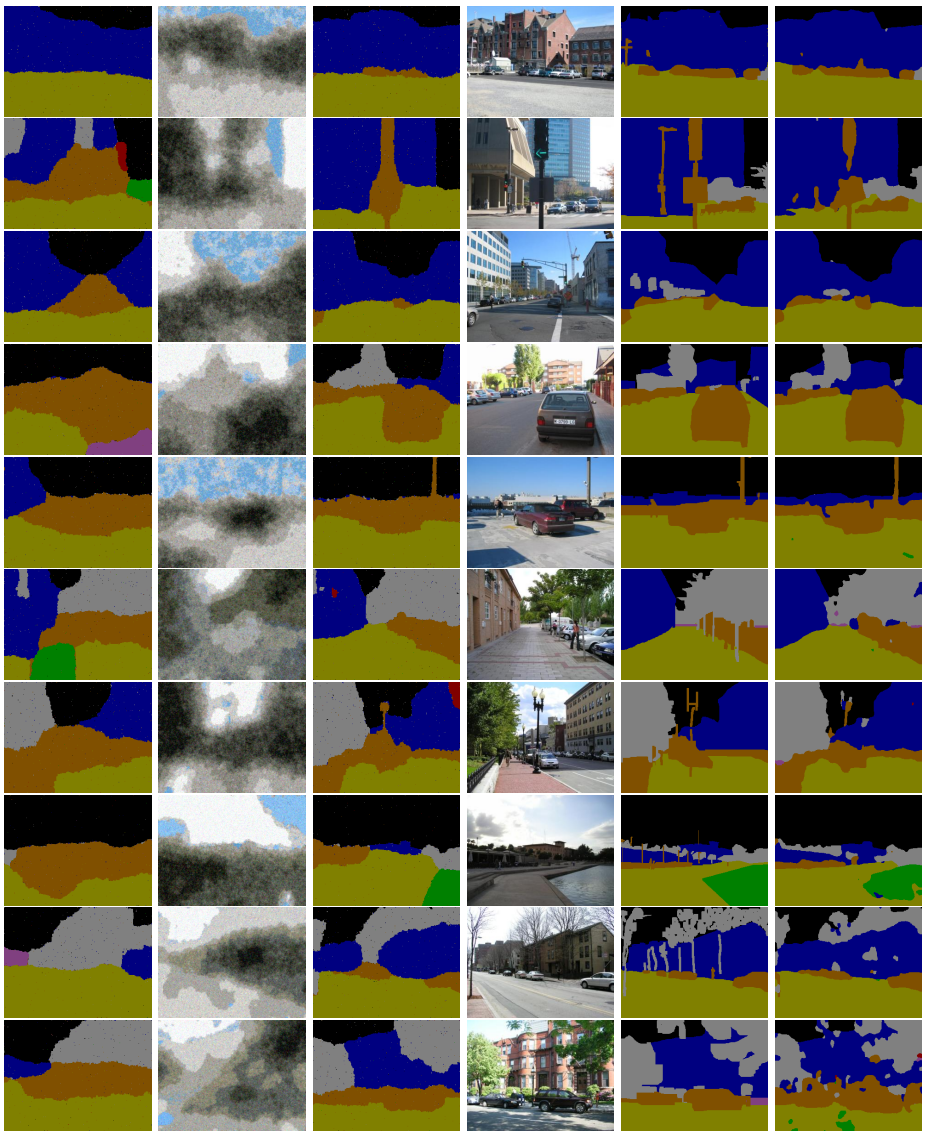

Figure 7: Chains of examples generated during the learning of implicit models. The first three columns are $\hat{y}$, $\tilde{x}$ and $\tilde{y}$ respectively (see Alg. 1). The fourth and fifth columns are samples from the training set, i.e. the observation x^* and the ground truth y^*. The last column is the max-marginal decision for the corresponding image.